

AD 646 719

RISK, AMBIGUITY, AND THE SAVAGE AXIOMS

Daniel Ellsberg

August 1961

BEST AVAILABLE COPY

RISK, AMBIGUITY, AND THE SAVAGE AXIOMS

Daniel Ellsberg*

The RAND Corporation, Santa Monica, California

I. ARE THERE UNCERTAINTIES THAT ARE NOT RISKS?

There has always been a good deal of skepticism about the behavioral significance of Frank Knight's distinction between "measurable uncertainty" or "risk," which may be represented by numerical probabilities, and "unmeasurable uncertainty" which cannot. Knight maintained that the latter "uncertainty" prevailed -- and hence that numerical probabilities were inapplicable -- in situations when the decision-maker was ignorant of the statistical frequencies of events relevant to his decision; or when a priori calculations were impossible; or when the relevant events were in some sense unique; or when an important, once-and-for-all decision was concerned.¹ (For this and subsequent footnotes, see end of paper.)

Yet the feeling has persisted that, even in these situations, people tend to behave "as if" they assigned numerical probabilities, or "degrees of belief," to the events impinging on their actions. However, it is hard either to confirm or to deny such a proposition in the absence of precisely-defined procedures for measuring these alleged "degrees of belief."

* Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

Research for this paper was done while the author was a member of the Society of Fellows, Harvard University, 1957. An earlier version was read before the Econometric Society at its December 1960 meeting in St. Louis; and the present version incorporating changes in Section III will appear in the November 1961 issue of the Quarterly Journal of Economics, in a symposium on decision-making under uncertainty, together with a contribution by William Fellner and a note on the present paper by Howard Raiffa. In revising Section III, the author was particularly stimulated by discussions with A. Madansky, T. Schelling, L. Shapley, and S. Winter.

What might it mean operationally, in terms of refutable predictions about observable phenomena, to say that someone behaves "as if" he assigned quantitative likelihoods to events: or to say that he does not? An intuitive answer may emerge if we consider an example proposed by Shackle, who takes an extreme form of the Knightian position that statistical information on frequencies within a large, repetitive class of events is strictly irrelevant to a decision whose outcome depends on a single trial. Shackle not only rejects numerical probabilities for representing the uncertainty in this situation; he maintains that in situations where all the potential outcomes seem "perfectly possible" in the sense that they would not violate accepted laws and thus cause "surprise," it is impossible to distinguish meaningfully (i.e., in terms of a person's behavior, or any other observations) between the relative "likelihoods" of these outcomes. In throwing a die, for instance, it would not surprise us at all if an ace came up on a single trial, nor if, on the other hand, some other number came up. So Shackle concludes:

Suppose the captains in a Test Match have agreed that instead of tossing a coin for a choice of innings they will decide the matter by this next throw of a die, and that if it shows an ace Australia shall bat first, if any other number, then England shall bat first. Can we now give any meaningful answer whatever to the question, "Who will bat first?" except "We do not know?"²

Most of us might think we could give better answers than that. We could say, "England will bat first," or more cautiously: "I think England will probably bat first." And if Shackle challenges us as to what we "mean" by that statement, it is quite natural to reply: "We'll bet on England; and we'll give you good odds."

It so happens that in this case statistical information (on the behavior of dice) is available and does seem relevant even to a "single shot" decision, our bet; it will affect the odds we offer. As Damon Runyon:

once said, "The race is not always to the swift nor the battle to the strong, but that's the way to bet." However, it is our bet itself, and not the reasoning and evidence that lies behind it, that gives operational meaning to our statement that we find one outcome "more likely" than another. And we may be willing to place bets -- thus revealing "degrees of belief" in a quantitative form -- about events for which there is no statistical information at all, or regarding which statistical information seems in principle unobtainable. If our pattern of bets were suitably orderly -- if it satisfied certain postulated constraints -- it would be possible to infer for ourselves numerical subjective probabilities for events, in terms of which some future decisions could be predicted or described. Thus a good deal -- perhaps all -- of Knight's class of "unmeasurable uncertainties" would have succumbed to measurement, and "risk" would prevail instead of "uncertainty."

A number of sets of constraints on choice-behavior under uncertainty have now been proposed, all more or less equivalent or closely similar in spirit, having the implication that -- for a "rational" man -- all uncertainties can be reduced to risks.³ Their flavor is suggested by Ramsay's early notions that, "The degree of a belief is...the extent to which we are prepared to act upon it," and "The probability of $1/3$ is clearly related to the kind of belief which would lead to a bet of 2 to 1."⁴ Starting from the notion that gambling choices are influenced by, or "reflect," differing degrees of belief, this approach sets out to infer those beliefs from the actual choices. Of course, in general those choices reveal not only the person's relative expectations but his relative preferences for outcomes; there is a problem of distinguishing between these.

but if one picks the right choices to observe, and if the Savage postulates or some equivalent set are found to be satisfied, this distinction can be made unambiguously, and either qualitative or, ideally, numerical probabilities can be determined. The propounders of these axioms tend to be hopeful that the rules will be commonly satisfied, at least roughly and most of the time, because they regard these postulates as normative maxims, widely-acceptable principles of rational behavior. In other words, people should tend to behave in the postulated fashion, because that is the way they would want to behave. At the least, these axioms are believed to predict certain choices that people will make when they take plenty of time to reflect over their decision, in the light of the postulates.

In considering only deliberate decisions, then, does this leave any room at all for "unmeasurable uncertainty": for uncertainties not reducible to "risks," to quantitative or qualitative probabilities?

A side effect of the axiomatic approach is that it supplies, at last (as Knight did not), a useful operational meaning to the proposition that people do not always assign, or act "as though" they assigned, probabilities to uncertain events. The meaning would be that with respect to certain events they did not obey, nor did they wish to obey -- even on reflection -- Savage's postulates or equivalent rules. One could emphasize here either that the postulates failed to be acceptable in those circumstances as normative rules, or that they failed to predict reflective choices; I tend to be more interested in the latter aspect, Savage no doubt in the former. (A third inference, which H. Raiffa favors, could be that people need more drill on the importance of conforming to the Savage axioms.) But from either point of view, it would follow that there would

be simply no way to infer meaningful probabilities for those events from their choices, and theories which purported to describe their uncertainty in terms of probabilities would be quite inapplicable in that area (unless quite different operations for measuring probability were devised). Moreover, such people could not be described as maximizing the mathematical expectation of utility on the basis of numerical probabilities for those events derived on any basis. Nor would it be possible to derive numerical "von Neumann-Morgenstern" utilities from their choices among gambles involving those events.

I propose to indicate a class of choice-situations in which many otherwise reasonable people neither wish nor tend to conform to the Savage postulates, nor to the other axiom sets that have been devised. But the implications of such a finding, if true, are not wholly destructive. First, both the predictive and normative use of the Savage or equivalent postulates might be improved by avoiding attempts to apply them in certain, specifiable circumstances where they do not seem acceptable. Second, we might hope that it is precisely in such circumstances that certain proposals for alternative decision rules and non-probabilistic descriptions of uncertainty (e.g., by Knight, Shackle, Hurvich, and Hodges and Lehmann) might prove fruitful. I believe, in fact, that this is the case.

FOOTNOTES

¹ Knight, F. H., Risk, Uncertainty and Profit, Houghton Mifflin Co., Boston, 1921. But see Arrow's comments: "In brief, Knight's uncertainties seem to have surprisingly many of the properties of ordinary probabilities, and it is not clear how much is gained by the distinctionActually, his uncertainties produce about the same reactions in individuals as other writers ascribe to risks." Arrow, K. J., "Alternative Approaches to the Theory of Choice in Risk-taking Situation," Econometrica, Vol. 19, October 1951, pp. 417, 426.

² Shackle, G. L. S., Uncertainty in Economics (Cambridge 1955), p. 8. If this example were not typical of a number of Shackle's work, it would seem almost unfair to cite it, since it appears so transparently inconsistent with commonly-observed behavior. Can Shackle really believe that an Australian captain who cared about batting first would be indifferent between staking this outcome on "heads" or on an ace?

³ Ramsey, F. P., "Truth and Probability" (1926) in The Foundations of Mathematics and Other Logical Essays, London, 1931; Savage, L. J., The Foundations of Statistics, New York, 1951; de Finetti, B., "Recent Suggestions for the Reconciliation of Theories of Probability," pp. 217-218 of Proceedings of the Second (1950) Berkeley Symposium on Mathematical

Statistics and Probability, Berkeley, 1951; Suppes, P. (see Suppes, P., Davidson, D., and Siegel, S., Decision-Making, Stanford, 1957). Closely Related approaches, in which individual choice behavior is presumed to be stochastic, have been developed by Luce, R. D., Individual Choice Behavior, New York, 1959, and Chipman, J. S., "Stochastic Choice and Subjective Probability," in Decisions, Values and Groups, ed. Willner, D., New York, 1960. Although the argument in this paper applies equally well to these latter stochastic axiom systems, they will not be discussed explicitly.

⁴ Ramsey, op. cit., p. 171.

⁵ Op. cit., p. 21. Savage notes that the principle, in the form of the rationale above, "cannot appropriately be accepted as a postulate in the sense that P1 is, because it would introduce new undefined technical terms referring to knowledge and possibility that would render it mathematically useless without still more postulates governing these terms." He substitutes for it a postulate corresponding to P2 above as expressing the same intuitive constraint. Savage's P2 corresponds closely to "Rubin's Postulate" (Luce and Raiffa, Games and Decisions, New York, 1957, p. 290) or Milnor's "Column Linearity" postulate, ibid., p. 297, which imply that adding a constant to a column of payoffs should not change the preference ordering among acts.

If numerical probabilities were assumed known, so that the subject were dealing explicitly with known "risks," these postulates would amount to Samuelson's "Special Independence Assumption" ("Probability, Utility, and the Independence Axiom," Econometrica, 20, 670-78, 1952) on which Samuelson relies heavily in his derivation of "von Neumann-Morgenstern utilities."

⁶ I bet.

⁷ Note that in no case are you invited to choose both a color and an urn freely; nor are you given any indication beforehand as to the full set of gambles that will be offered. If these conditions were altered (as in some of H. Raiffa's experiments with students), you could employ randomized strategies, such as flipping a coin to determine what color to bet on in Urn I, which might affect your choices.

⁸ Here we see the advantages of purely hypothetical experiments. In "real life," you would probably turn out to have a profound color preference that would invalidate the whole first set of trials, and various other biases that would show up one by one as the experimentation progressed inconclusively.

However, the results in Chipman's almost identical experiment (op. cit., pp. 87-88) do give strong support to this finding; Chipman's explanatory hypothesis differs from that proposed below.

⁹ In order to relate these choices clearly to the postulates, let us change the experimental setting slightly. Let us assume that the balls in Urn I are each marked with a I, and the balls in Urn II with a II; the contents of both urns are then dumped into a single urn, which then contains 50 Red_{II} balls, 50 Black_{II} balls, and 100 Red_I and Black_I balls in unknown proportion (or in a proportion indicated only by a small random sample, say, one Red and one Black). The following actions are to be considered:

	100		50	50
	R_I	B_I	R_{II}	B_{II}
I	a	b	b	b
II	b	a	b	b
III	b	b	a	b
IV	b	b	b	a
V	a	a	b	b
VI	b	b	a	a

Let us assume that a person is indifferent between I and II (between betting on R_I or B_I), between III and IV and between V and VI. It would then follow from Postulates 1 and 2, the assumption of a complete ordering of actions and the Sure-thing Principle, that I, II, III and IV are all indifferent to each other.

To indicate the nature of the proof, suppose that I is preferred to III (the person prefers to bet on R_I rather than R_{II}). Postulates 1 and 2 imply that certain transformations can be performed on this pair of actions without affecting their preference ordering; specifically, one action can be replaced by an action indifferent to it (P1 -- complete ordering) and the value of a constant column can be changed (P2 -- Sure-thing Principle)

Thus starting with I and III and performing such "admissible transformations" it would follow from P1 and P2 that the first action in each of the following pairs should be preferred:

	R_I	R_{II}	R_{III}	R_{IV}	
I	a	b	b	b	
III	b	b	a	b	
<hr/>					
I'	a	b	b	a	P2
III'	b	b	a	a	
<hr/>					
I''	a	b	b	a	P1
III''	a	a	b	b	
<hr/>					
I'''	b	b	b	a	P2
III'''	b	a	b	b	
<hr/>					
I''''	b	b	a	b	P1
III''''	a	b	b	b	

Contradiction: I preferred to III, and I'''' (equivalent to III) preferred to III'''' (equivalent to I).

¹⁰ Knight, *op. cit.*, p. 219.

¹¹ Kenneth Arrow has suggested the following example, in the spirit of the above one:

	100		50	50
	R_I	R_{II}	R_{III}	R_{IV}
I	a	a	b	b
II	a	b	a	b
III	b	a	b	a
IV	b	b	a	a

Assume that I is indifferent to IV, II is indifferent to III.

Suppose that I is preferred to II; what is the ordering of III and IV? If III is not preferred to IV, P2, the Sure-thing Principle is violated. If IV is not preferred to III, P1, complete ordering of actions, is violated. (If III is indifferent to IV, both P1 and P2 are violated.)

¹² Let the utility payoffs corresponding to \$100 and \$0 be 1, 0; let P_1, P_2, P_3 be the probabilities corresponding to Red, Yellow, Black. The expected value to action I is then P_1 ; to II, P_2 ; to III, $P_1 + P_3$; to

IV, $P_2 + P_3$. But there are no P's, $P_1 \geq 0$, $\Sigma P_i = 1$, such that $P_1 > P_2$ and $P_1 + P_3 < P_2 + P_3$.

¹³ Samuelson, P., "Probability and the Attempts to Measure Utility," The Economic Review (Tokyo, Japan), July 1950, pp. 169-170.

To test the predictive effectiveness of the axioms (or of the alternate decision rule to be proposed in the next section) in these situations controlled experimentation is in order. (See Chipman's ingenious experiment op. cit.) But, as Savage remarks (op. cit., p. 26), the mode of interrogation implied here and in Savage's book, asking "the person not how he feels but what he would do in such and such a situation" and giving him ample opportunity to ponder the implications of his replies, seems quite appropriate in weighing "the theory's more important normative interpretation." Moreover, these non-experimental observations can have at least negative empirical implications, since there is a presumption that people whose instinctive choices violate the Savage axioms, and who claim upon further reflection that they do not want to obey them, do not tend to obey them usually in such situations.

¹⁴ No one whose decisions were based on "regrets" could violate the Sure-thing Principle, since all constant columns of payoffs would transform to a column of 0's in terms of "regret"; on the other hand, such a person would violate P1, complete ordering of strategies.

¹⁵ See Chipman, op. cit., pp. 75, 93. Chipman's important work in this area, done independently and largely prior to mine, is not discussed here since it embodies a stochastic theory of choice; its spirit is otherwise closely similar to that of the present approach, and his experimental results are both pertinent and favorable to the hypotheses below

(though Chipman's inferences are somewhat different).

See also the comments by N. Georgescu-Roegen on notion of "credibility," a concept identical to "ambiguity" in this paper: "The Nature of Expectation and Uncertainty," in Expectations, Uncertainty, and Business Behavior, ed. Mary Brown, Social Science Research Council, New York, 1958, pp. 24-26; and "Choice, Expectations and Measurability," Quarterly Journal of Economics, Vol. LXVIII, No. 4, November 1954, pp. 527-530. These highly pertinent articles came to my attention only after this paper had gone to the printer, allowing no space for comment here.

¹⁶ Savage, op. cit., pp. 57-58, 59. Savage later goes so far as to suggest (op. cit., pp. 168-169) that the "aura of vagueness" attached to many judgments of personal probability might lead to systematic violations of his axioms, although the decision rule he discusses as alternative--minimizing regret--cannot, as mentioned in footnote 14 above, account for the behavior in our examples.

¹⁷ Knight, op. cit., p. 227.

¹⁸ This contradicts the assertions by Chipman (op. cit., p. 28) and Georgescu-Roegen ("Choice, Expectations and Measurability," pp. 527-530), and "The Nature of Expectation and Uncertainty," p. 25) that individuals order uncertainty-situations lexicographically in terms of estimated expectation and "credibility" (ambiguity); ambiguity appears to influence choice even when estimated expectations are not equivalent.

¹⁹ This rule is based upon the concept of a "restricted Bayes solution" developed by J. L. Hodges, Jr., and E. L. Lehmann ("The Uses of Previous Experience in Reaching Statistical Decision," Annals of

Mathematical Statistics, Vol. 23, No. 3, September, 1952, pp. 396-407. The discussion throughout Section III of this paper derives heavily from the Hodges and Lehmann argument, although their approach is motivated and rationalized somewhat differently.

See also, L. Hurvics, "Some Specification Problems and Applications to Econometric Models," Econometrica, Vol. 19, No. 3, July, 1951, pp. 343-344 (abstract). This deals with the same sort of problem and presents a "generalized Bayes-minimax principle" equivalent, in more general form, to the decision rule I proposed in an earlier presentation of this paper (December, 1960); but both of these lacked the crucial notions developed in the Hodges and Lehmann approach of a "best estimate" distribution y^0 and a "confidence" parameter ρ .

²⁰ This interpretation of the behavior-pattern contrasts to the hypothesis or decision rule advanced by Fellner in the accompanying article in this symposium. Fellner seems unmistakably to be dealing with the same phenomena discussed here, and his proposed technique of measuring a person's subjective probabilities and utilities in relatively "unambiguous" situations and then using these measurements to calibrate his uncertainty in more ambiguous environments seems to me a most valuable source of new data and hypotheses. Moreover, his descriptive data and intuitive conjectures lend encouraging support to the findings reported here. However, his solution to the problem supposes a single set of weights determined independently of payoffs (presumably corresponding to the "best estimates" here) and a "correction factor," reflecting the degree of ambiguity or confidence, which operates on these weights in a manner independent of the structure of payoffs. I am not entirely clear

on the behavioral implications of Fellner's model or the decision rule it implies, but in view of these properties I am doubtful whether it can account adequately for all the behavior discussed above.